



IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

AF
2143

Applicant(s): LANCE W. RUSSELL

Confirmation No.: 8674

Application No.: 09/895,235

Examiner: Bilgrami, Asghar

Filing Date: June 28, 2001

Group Art Unit: 2143

Title: MIGRATING RECOVERY MODULES IN A DISTRIBUTED COMPUTING ENVIRONMENT

Mail Stop Appeal Brief-Patents
Commissioner For Patents
PO Box 1450
Alexandria, VA 22313-1450

TRANSMITTAL OF APPEAL BRIEF

Sir:

Transmitted herewith is the Appeal Brief in this application with respect to the Notice of Appeal filed on Aug. 9, 2005.

The fee for filing this Appeal Brief is (37 CFR 1.17(c)) \$500.00.

(complete (a) or (b) as applicable)

The proceedings herein are for a patent application and the provisions of 37 CFR 1.136(a) apply.

() (a) Applicant petitions for an extension of time under 37 CFR 1.136 (fees: 37 CFR 1.17(a)-(d) for the total number of months checked below:

- | | |
|------------------|-----------|
| () one month | \$120.00 |
| () two months | \$450.00 |
| () three months | \$1020.00 |
| () four months | \$1590.00 |

() The extension fee has already been filled in this application.

(X) (b) Applicant believes that no extension of time is required. However, this conditional petition is being made to provide for the possibility that applicant has inadvertently overlooked the need for a petition and fee for extension of time.

Please charge to Deposit Account **08-2025** the sum of **\$500.00**. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16 through 1.21 inclusive, and any other sections in Title 37 of the Code of Federal Regulations that may regulate fees. A duplicate copy of this sheet is enclosed.

(X) I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to:
Commissioner for Patents, Alexandria, VA
22313-1450. Date of Deposit: Oct. 10, 2005
OR

() I hereby certify that this paper is being transmitted to the Patent and Trademark Office facsimile number _____ on _____

Number of pages:

Typed Name: **Edouard Garcia**

Signature:

Respectfully submitted,

LANCE W. RUSSELL

By

Edouard Garcia

Attorney/Agent for Applicant(s)
Reg. No. **38,461**

Date: **Oct. 10, 2005**

Telephone No.: **(650) 289-0904**



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant : LANCE W. RUSSELL Art Unit : 2143
Serial No. : 09/895,235 Examiner : Bilgrami, Asghar H.
Filed : June 28, 2001
Title : MIGRATING RECOVERY MODULES IN A DISTRIBUTED COMPUTING ENVIRONMENT

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEFI. Real Party in Interest

The real party in interest is Hewlett-Packard Development Company, L.P., a Texas Limited Partnership having its principal place of business in Houston, Texas.

II. Related Appeals and Interferences

Appellant is not aware of any related appeals or interferences that will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

III. Status of Claims

Claims 1-25, which are the subject of this appeal, are pending.

Claims 1-25 stand rejected.

Appellant appeals all rejections of the pending claims 1-25.

10/14/2005 EFLORES 00000096 082025 09895235

01 FC:1402 500.00 DA

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as First Class Mail in an envelope addressed to: Commission for Patents, PO Box 1450, Alexandria, VA 22313-1450 on:

October 10, 2005

Date

(Signature of person mailing papers)

Edouard Garcia

(Typed or printed name of person mailing papers)

Applicant : Lance W. Russell
Serial No. : 09/895,235
Filed : June 28, 2001
Page : 2 of 16

Attorney's Docket No.: 10003532-1
Appeal Brief dated Oct. 10, 2005
Reply to final action dated May 12, 2005

IV. Status of Amendments

The sole Amendment filed December 23, 2004, has been entered and acted upon by the Examiner.

No amendments were filed after the final Office action dated May 12, 2005.

V. Summary of Claimed Subject Matter

The invention claimed in independent claim 1 is a system for managing a plurality of distributed nodes of a network. The system includes a recovery module that migrates from one network node to another, determines a status of a network node, and initiates a recovery process on a failed network node.

FIG. 1 shows a network management node 12 that includes a network management module that is configured to manage a plurality of distributed nodes 14, 16 of a distributed computing system 10 in accordance with the invention defined in independent claim 1. FIG. 1 shows the network management node 12 launching a migratory recovery module 20 into the distributed computing system 10. The Specification explains that the recovery module 20 migrates from node to node, determines the status of each node, and initiates recovery processes on failed nodes (see, e.g., page 5, lines 13-15). FIG. 2 shows an embodiment of the network management node 12. FIG. 3 shows an embodiment of a method implemented by the network management module for monitoring the status and implementing recovery processes on one or more nodes of the distributed computer system 10.

The invention recited in independent claim 11 is a method for managing a plurality of distributed nodes of a network. In accordance with this method, a status of a current one of the network nodes is determined. In response to a determination that the current network node has failed, a recovery process is initiated on the current network node. After migrating from the current network node to a successive one of the network nodes, the process is repeated.

FIG. 5 shows an embodiment of a method that is executed by the recovery module 20 in accordance with the invention defined in independent claim 11. In accordance with this method, the recovery module 20 monitors the health of a network node (FIG. 5, step 80; page 9, line 27 through page 10, line 3). If the recovery module 20 determines that one or more

node processes have failed (FIG. 5, step 82), the recovery module 20 initiates a recovery process on the network node in accordance with a restart protocol (FIG. 5, step 84; page 10, lines 3-7). After reporting the node status to the network management module (FIG. 5, step 86), the recovery module 20 requests to be transmitted to the next destination node (FIG. 5, step 88; page 10, lines 10-20).

The Specification explains that the recovery module 20 may include a routing method that may be invoked to identify the destination node address based on a routing table stored at the current network node in accordance with the invention defined in claims 2, 3, 12, and 13 (see, e.g., page 10, lines 10-18).

VI. Grounds of Rejection to be Reviewed on Appeal

Claims 1-25 stand rejected under 35 U.S.C § 102(e) over Forbes (U.S. 6,728,896).

VII. Argument

A. Rejection under 35 U.S.C. § 102(e) over Forbes (U.S. 6,728,896)

1. Claims 1, 4-10, and 21-25

Independent claim 1 recites that the system includes “a recovery module configured to migrate from one network node to another, determine a status of a network node, and initiate a recovery process on a failed network node.”

In the final Office action dated May 12, 2005, the Examiner has stated that:

As per claims 1, 10, 11, 19, 20, 21-24 & 25 Forbes disclosed a method for managing a plurality of distributed nodes of a network, comprising: (a) on a current one of the network nodes, determining a status of the current [[a]] network node. Forbes describes “heart beats” as means for detecting the status of the nodes (col. 4, lines 27-33 & col. 9, lines 1-18); [[and]] (b) in response to a determination that the current network node has failed, initiating a recovery process on the current network node; (c) migrating from the current network node to a successive one of the network nodes; and (d) repeating (a), (b), and (c) with the current node corresponding to the successive

network. (col. 3, lines 48-67, col. 4, lines 1-33, 53-65 & col. 11, lines 15-23).

The sections of Forbes' disclosure in columns 3 and 4 that the Examiner has cited in support of his rejection of claim 1 describe the heart beat monitoring and failure recovery initiation functions of the Microsoft Cluster Server (MSCS) software that is described in Forbes. Therefore, the Examiner's position appears to be that the MSCS corresponds to the recovery module that is recited in claim 1.¹

As explained in detail below, however, in accordance with Forbes' disclosure the MSCS does not migrate from one network node to another. Indeed, contrary to the Examiner's statement quoted above, Forbes does not disclose anything that corresponds to a recovery module that is configured to migrate from one network node to another, determine a status of a network node, and initiate a recovery process on a failed network node, as recited in independent claim 1.

Forbes discloses a failover method of a simulated operating system in a clustered computing environment. In accordance with Forbes' approach, first and second servers have respective copies of first and second operating systems (see, e.g., col. 3, lines 4-30, and FIG. 3). Initially, the first server executes its copies of both the first and second operating systems, whereas the second server executes only its copy of the first operating system (see, e.g., col. 3, lines 4-120). In response to the detection of the failure of the first server, the second server executes its copy of the second operating system (see, e.g., col. 3, lines 4-12). In this way, only a single copy of the second operating system is executing on the first and second servers at a time, in compliance with a single-server license (see, e.g., col. 3, lines 25-29). The detection of failure is done by sensing a "heartbeat" communication on a "Private Network" connection between the first and second servers, where the disappearance of the heartbeat signal from a server indicates a failure of that server (see, e.g., col. 3, lines 18-24).

In Forbes' preferred embodiment, the first operating system is a Windows operating system, such as Windows NT or Windows 2000 (see, e.g., col. 10, lines 47-55) and the second operating system is the Unisys MCP mainframe operating system (Master Control Program) (see, e.g., col. 10, lines 56-60). An adapting means, such as the ClearPath Virtual

¹ The section of Forbes' disclosure in column 11 that the Examiner has cited relates to the function of the adapting means that allows a second operating system (i.e., the MCP operating system) to operate on top of the Windows operating system. The adapting means does not perform heartbeat monitoring or failure recovery initiation functions.

Machine, allows the MCP operating system to execute on top of the Windows operating system (see, e.g., col. 11, lines 1-3). The adapting means is integrated with Microsoft Cluster Services (MSCS), which supports the connection of the first and second servers into a cluster (see col. 2, lines 43-52, and col. 3, lines 48-51). Each of the first and second servers executes its own copy of the adapting means (see, e.g., FIG. 3).

The Examiner has not pointed to any disclosure in Forbes that teaches that the MSCS migrates from one server to another. This is not surprising since Forbes does not teach or suggest anything that would have led one of ordinary skill in the art at the time the invention was made to believe that the MSCS migrates from one node to another. To the contrary, Forbes expressly teaches that each of the first and second servers executes its own copy of the MSCS. For example, FIG. 2 shows that the MSCS is a part of the Windows operating system (see col. 5, lines 23-28 and col. 7, lines 19-25) and FIG. 3 shows that each of the first and second servers executes its own copy of the Windows operating system. Thus, a respective copy of the MSCS executes on each of the first and second servers along with a respective copy of the Windows operating system. Indeed, if each server did not execute a respective copy of the MSCS, the servers would not be able to form a cluster. Since each server is executing a respective copy of the MSCS, there is no need whatsoever for the MSCS to migrate from one server to another.

In response to Applicant's indication that the Examiner did not point specifically to the component of Forbes' system that he believed corresponded to the recovery module that is recited in claim 1 (see page 7 of the Amendment dated December 23, 2004), the Examiner has stated that (page 4, ¶ 14 of the final Office action dated May 12, 2005):

As to applicants arguments Forbes disclosed the recovery module (MSCS) which can detect and recover from server or application failures (see col. 3, lines 48-67) and also talks about recovering its previous state after being switched over to a second server once the first server failed on the network (col. 11, lines 15-22).

In this statement, however, the section of Forbes' disclosure that the Examiner has cited to support his incorrect contention that Forbes discloses that the MSCS is "switched over to a second server" relates only to the adapting means (i.e., the Virtual Machine for ClearPath MCP software) and the second operating system (i.e., the MCP operating system); this disclosure does not relate to the MSCS. Indeed, the entire point of Forbes' invention is to

enable the MCP operating system to execute on the surviving server only after the other server has failed so that only a single copy of the second operating system is executing on the first and second servers at a time, in compliance with a single-server license (see, e.g., col. 3, lines 25-29).

It is noted that the heartbeat message that is transmitted from one server to another also does not constitute the recovery module that is recited in claim 1. In particular, the heartbeat messages are not configured to determine a status of a server. The heartbeat messages merely indicate the status of the transmitting servers (e.g., the status of the first server is determined by the MSCS operating on the second server, which expects to receive a heartbeat message from the first server). The heartbeat messages also are not configured to initiate a recovery process on a failed server. Instead, the MSCS operating on a surviving server transfers the ownership of resources (e.g., disk drives and IP addresses) from a failed server to the surviving server and restarts the workload of the failed server on the surviving server (see col. 9, lines 8-18 and 46-51).

In summary, Forbes does not disclose “a recovery module configured to migrate from one network node to another, determine a status of a network node, and initiate a recovery process on a failed network node,” as recited in claim 1. Accordingly, the Examiner’s rejection of independent claim 1 under 35 U.S.C. § 102(e) over Forbes should be withdrawn.

Each of claims 4-10 and 21-25 incorporates the features of independent claim 1 and therefore is patentable over Forbes for at least the same reasons explained above.

2. Claims 2 and 3

Each of claims 2 and 3 incorporates the features of independent claim 1 and therefore is patentable over Forbes for at least the same reasons explained above. Claims 2 and 3 also are patentable for the following additional reasons.

Claim 2 recites that “the recovery module comprises a routing component for determining next hop addresses for migrating the recovery module from an origin network node to a series of successive destination network nodes.”

In his rejection of claim 2, the Examiner has asserted that in col. 9, lines 8-12, Forbes discloses a routing component for determining a next hop address from an origin network

node to a destination network node (see page 2, ¶ 4 of the final Office action dated May 12, 2005). The disclosure in col. 9, lines 8-12, however, recites:

In the event of a server failure, the MSCS software employs a “shared nothing” clustering architecture that automatically transfers ownership of resources (such as disk drives and IP addresses) from a failed server over to a surviving server.

The resources whose ownership is transferred to the surviving server in Forbes' approach, however, do not have anything to do with next hop addresses for migrating a recovery module from an origin network node to a series of successive destination network nodes, as recited in claim 2. The IP addresses in particular are merely the IP addresses that were assigned to the failed server, not next hop addresses. Moreover, as explained above in connection with independent claim 1, neither the MSCS nor the heartbeat messages migrate from an origin network node to a series of successive destination network nodes. Therefore, it would not have served any useful purpose for the MSCS or the heartbeat messages to incorporate “a routing component for determining next hop addresses for migrating the recovery module from an origin network node to a series of successive destination network nodes,” as recited in claim 2.

Claim 3 incorporates the features of claim 2 and therefore is patentable over Forbes for at least the same reasons.

3. Independent claim 11

Claim 11 recites:

11. A method for managing a plurality of distributed nodes of a network, comprising:
 - (a) on a current one of the network nodes, determining a status of the current network node;
 - (b) in response to a determination that the current network has failed, initiating a recovery process on the current network node;
 - (c) migrating from the current network node to a successive one of the network nodes; and
 - (d) repeating (a), (b), and (c) with the current network node corresponding to the successive network node for each of the nodes in the network.

The Examiner has rejected independent claim 11 on the same basis as independent claim 1 (see page 2, ¶ 3 of the final Office action dated May 12, 2005). As explained above in connection with claim 1, however, Forbes does not disclose anything that is capable of determining a status of the current network node, initiating a recovery process on the current network node, and migrating from the current network node to a successive network node, as recited in claim 11.

For at least these reasons, the Examiner's rejection of independent claim 11 under 35 U.S.C. § 102(e) over Forbes now should be withdrawn.

Each of claims 14-19 incorporates the features of independent claim 11 and there is patentable over Forbes for at least the same reasons explained above.

4. Claims 12 and 13

Each of claims 12 and 13 incorporates the features of independent claim 11 and therefore is patentable over Forbes for at least the same reasons explained above. Claims 12 and 13 also are patentable for the following additional reasons.

Claim 12 recites "migrating from one network node to another comprises determining a next hop address from an origin network node to a destination network node."

In his rejection of claim 12, the Examiner has asserted that in col. 9, lines 8-12, Forbes discloses a routing component for determining a next hop address from an origin network node to a destination network node (see page 2, ¶ 4 of the final Office action dated May 12, 2005). The disclosure in col. 9, lines 8-12, however, recites that:

In the event of a server failure, the MSCS software employs a "shared nothing" clustering architecture that automatically transfers ownership of resources (such as disk drives and IP addresses) from a failed server over to a surviving server.

The resources whose ownership is transferred to the surviving server in Forbes' approach, however, do not have anything to do with next hop addresses for migrating a recovery module from an origin network node to a series of successive destination network nodes, as recited in claim 12. The IP addresses in particular are merely the IP addresses that were assigned to the failed server, not next hop addresses. Moreover, as explained above in connection with independent claim 1, neither the MSCS nor the heartbeat messages migrate

from an origin network node to a series of successive destination network nodes. Therefore, it would not have served any useful purpose for the MSCS or the heartbeat messages to be capable of “migrating from one network node to another comprises determining a next hop address from an origin network node to a destination network node,” as recited in claim 12.

Claim 13 incorporates the features of claim 12 and therefore is patentable over Forbes for at least the same reasons.

5. Independent claim 20

Claim 20 recites that the computer program comprises computer-readable instructions for causing a computer to:

migrate the computer program from one network node to a series of successive network nodes;

determine a status of each network node to which the computer program is migrated; and

initiate a recovery process on each network node to which the computer program is migrated and is determined to have failed.

The Examiner has rejected independent claim 20 on the same basis as independent claim 1 (see page 2, ¶ 3 of the final Office action dated May 12, 2005). As explained above in connection with claim 1, however, Forbes does not disclose anything that corresponds to computer-readable instructions for causing a computer to determine a status of the current network node, initiate a recovery process on the current network node, and migrate from the current network node to a successive network node, as recited in claim 20.

For at least these reasons, the Examiner’s rejection of independent claim 20 under 35 U.S.C. § 102(e) over Forbes now should be withdrawn.

VIII. Conclusion

For the reasons explained above, all of the pending claims are now in condition for allowance and should be allowed.

Charge any excess fees or apply any credits to Deposit Account No. 08-2025.

Applicant : Lance W. Russell
Serial No. : 09/895,235
Filed : June 28, 2001
Page : 10 of 16

Attorney's Docket No.: 10003532-1
Appeal Brief dated Oct. 10, 2005
Reply to final action dated May 12, 2005

Respectfully submitted,



Edouard Garcia
Reg. No. 38,461
Telephone No.: (650) 631-6591

Please direct all correspondence to:

Hewlett-Packard Company
Intellectual Property Administration
Legal Department, M/S 35
P.O. Box 272400
Fort Collins, CO 80528-9599

CLAIMS APPENDIX

The claims that are the subject of Appeal are presented below.

Claim 1 (original): A system for managing a plurality of distributed nodes of a network, comprising:

a recovery module configured to migrate from one network node to another, determine a status of a network node, and initiate a recovery process on a failed network node.

Claim 2 (previously presented): The system of claim 1, wherein the recovery module comprises a routing component for determining next hop addresses for migrating the recovery module from an origin network node to a series of successive destination network nodes.

Claim 3 (previously presented): The system of claim 2, wherein the routing component is configured to determine the next hop addresses based upon a routing table stored at the origin network node.

Claim 4 (original): The system of claim 1, wherein the recovery module is configured to determine the status of a network node by sending an inter-process communication to a node process.

Claim 5 (original): The system of claim 1, wherein the recovery module is configured to determine the status of a network node in accordance with a heartbeat messaging protocol.

Claim 6 (original): The system of claim 1, wherein the recovery module is configured to initiate a recovery process on a failed network node in accordance with a restart protocol.

Claim 7 (original): The system of claim 6, wherein the recovery module is configured to initiate a restart of a failed node process by transmitting a request to a process execution service operating on the failed network node.

Claim 8 (original): The system of claim 1, wherein the recovery module is configured to transmit a node status message to a network management module operating at a network management network node.

Claim 9 (original): The system of claim 8, wherein the node status message comprises information obtained from a log file generated at the failed network node.

Claim 10 (original): The system of claim 1, further comprising a network management module configured to launch a plurality of recovery modules into the network.

Claim 11 (previously presented): A method for managing a plurality of distributed nodes of a network, comprising:

- (a) on a current one of the network nodes, determining a status of the current network node;
- (b) in response to a determination that the current network has failed, initiating a recovery process on the current network node;
- (c) migrating from the current network node to a successive one of the network nodes; and
- (d) repeating (a), (b), and (c) with the current network node corresponding to the successive network node for each of the nodes in the network.

Claim 12 (original): The method of claim 11, wherein migrating from one network node to another comprises determining a next hop address from an origin network node to a destination network node.

Claim 13 (original): The method of claim 12, wherein the next hop address is determined based upon a routing table stored at the origin network node.

Claim 14 (original): The method of claim 11, wherein the status of a network node is determined by sending an inter-process communication to a node process.

Claim 15 (original): The method of claim 11, wherein the status of a network node is determined in accordance with a heartbeat messaging protocol.

Claim 16 (original): The method of claim 11, wherein a recovery process is initiated on a failed network node in accordance with a restart protocol.

Claim 17 (original): The method of claim 16, wherein a restart of a failed node process is initiated by transmitting a request to a process execution service operating on the failed network node.

Claim 18 (original): The method of claim 11, further comprising transmitting a node status message to a network management module operating at a network management network node.

Claim 19 (original): The method of claim 11, further comprising launching into the network a plurality of recovery modules, each configured to migrate from one network node to another, determine the status of a network node, and initiate a recovery process on a failed network node.

Claim 20 (previously presented): A computer program for managing a plurality of distributed nodes of a network, the computer program residing on a computer-readable medium and comprising computer-readable instructions for causing a computer to:

migrate the computer program from one network node to a series of successive network nodes;

determine a status of each network node to which the computer program is migrated; and

initiate a recovery process on each network node to which the computer program is migrated and is determined to have failed.

Claim 21 (previously presented): The system of claim 1, wherein the recovery module is a software object that is instantiatable by a respective operating environment on each network node.

Claim 22 (previously presented): The system of claim 21, wherein the operating environment on each of the network nodes provides the recovery module with access to status monitoring resources, recovery resources, and native operative system resources that are available at each of the network nodes.

Claim 23 (previously presented): The system of claim 1, wherein, upon migrating from a first one of the network nodes to a second one of the network nodes and being instantiated on the second node, the recovery module determines a status of the second network node.

Claim 24 (previously presented): The system of claim 23, wherein the recovery module initiates the recovery process on the second node in response to a determination that the second node has failed.

Claim 25 (previously presented): The system of claim 23, wherein the recovery module is configured to migrate to a third one of the network nodes after determining the status of the second network node.

Applicant : Lance W. Russell
Serial No. : 09/895,235
Filed : June 28, 2001
Page : 15 of 16

Attorney's Docket No.: 10003532-1
Appeal Brief dated Oct. 10, 2005
Reply to final action dated May 12, 2005

EVIDENCE APPENDIX

There is no evidence submitted pursuant to 37 CFR §§ 1.130, 1.131, or 1.132 or any other evidence entered by the Examiner and relied upon by Appellant in the pending appeal. Therefore, no copies are required under 37 CFR § 41.37(c)(1)(ix) in the pending appeal.

Applicant : Lance W. Russell
Serial No. : 09/895,235
Filed : June 28, 2001
Page : 16 of 16

Attorney's Docket No.: 10003532-1
Appeal Brief dated Oct. 10, 2005
Reply to final action dated May 12, 2005

RELATED PROCEEDINGS APPENDIX

Appellant is not aware of any decisions rendered by a court or the Board that will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal. Therefore, no copies are required under 37 CFR § 41.37(c)(1)(x) in the pending appeal.